Published in final edited form as:

Transp Res Part F Traffic Psychol Behav. 2019 April ; 62: 316–326. doi:10.1016/j.trf.2019.01.010.

Learning to drive: A reconceptualization

Jessica Hafetz Mirman^{a,*}, Allison E. Curry^b, and Daniel Mirman^a

^aThe University of Alabama at Birmingham, Department of Psychology, 1300 University Blvd, Birmingham, AL 35294

^bPerelman School of Medicine, University of Pennsylvania, Department of Pediatrics, 2716 South Street, 13th Floor, Philadelphia, PA 19146

Abstract

Drivers' population-level crash rates incrementally decrease following licensure, which has led to the implicit assumption that an individual driver's crash risk also decreases incrementally after licensure as they accrue experience. However, in the aggregate data an incremental decrease in crash rate can reflect both incremental reductions in crash risk within individuals and an incremental increase in the proportion of drivers who have experienced an abrupt decrease in crash risk. Therefore, while it is true to say that the population of drivers' crash risk reduces in the months following licensure, it is not necessarily true to say that a driver's crash risk reduces in the months following licensure; that is, it cannot be assumed that individual-level changes in crash risk mirror the population-level changes in crash rates. In statistics, this is known as an ecological fallacy and in formal logic it is known as the fallacy of division, a type of category error. Using computational cognitive modeling methods we demonstrate that aggregating individual-level abrupt decreases in crash risk (i.e., non-incremental change trajectories) accurately fits population-level crash rate data from over 1 million adolescents *and* uniquely accounts for effects of two interventions found to reduce police-reported MVCs. Thus, we demonstrate that (1) a power-law artifact is readily observable in newly licensed drivers' aggregate crash data, which is not necessarily indicative of individual-level change processes, (2) interventions can alter crash risk trajectories by inducing immediate phase changes in crash risk into a lower risk stratum, or increasing the probability of such a change, (3) a phase transition model provides a stronger and more parsimonious account of the existing data than an incremental-accrual model.

Keywords

strategic learning; phase transitions; teen driver safety; ecological fallacy; fallacy of division

*Corresponding Author 205.975.3907, 205.975.6110 (fax), jhmirman@uab.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1.0 Introduction

Prior studies of drivers' population-level crash data from jurisdictions with Graduated Driver Licensing (GDL) programs have shown that motor vehicle crash (MVC) rates peak at the transition to independent licensure, following a supervised learner's permit period, and then decrease incrementally with the most rapid reductions occurring during the first few months of licensure (Curry, Pfeiffer, Durbin, & Elliott, 2015; Foss, Martell, Goodwin, & O'Brien, 2011; Mayhew, Simpson, & Pak, 2003; McCartt, Mayhew, Braitman, Ferguson, & Simpson, 2009; Simons-Morton & Ehsani, 2016). It has been assumed that the incremental decrease in population-level crash rates means that individual drivers *also* have incrementally decreasing crash risk (McCartt et al., 2009; Simons-Morton & Ehsani, 2016). In a comprehensive analysis of population-level crash patterns McCartt et al., 2009 state, *"there remains no proven effective alternative to the learning that occurs incrementally when unsupervised driving begins. Although the largest gains from experience occur soon after licensure, the benefits continue to accrue even after several additional years.* (pg. 218) "(McCartt et al., 2009).

The implicit assumption that an individual driver's crash risk trajectory (i.e., the learning trajectory) decreases incrementally is directly inspired from patterns in the population-level, post-license crash data that resemble a conventional "learning curve" (Foss et al., 2011; McCartt et al., 2009; Simons-Morton & Ehsani, 2016). Indeed, this is a highly robust and visually evocative pattern in the population-level data that is frequently referred to as the "learning curve for driving" (Simons-Morton & Ehsani, 2016; IOM & NRC, 2011). However, it cannot be assumed that an individual driver's crash risk trajectory follows the population-level crash rate pattern, or put another way, that the crash risk trajectory of any given driver, or the "average driver", will also incrementally decrease as he or she incrementally accrues experience.

In the case of human learning and performance many studies have shown that, when aggregated, individual trajectories of *discontinuous* change patterns can take the appearance of an incremental change pattern (Gray & Lindstedt, 2017; Haider & Frensch, 2002; Murre & Chessa, 2011; Seigler, 2006). Indeed, Myung et al., 2000 demonstrated that a power-law "artifact" can occur when data from non-linear models are arithmetically averaged in the presence of individual difference factors (Myung, Kim, & Pitt, 2000). A conceptual illustration of this phenomenon applied to drivers' crash risk is shown in Figure 1.

The left two panels in Figure 1 contain simulated crash risk curves for 30 individual drivers based on an incremental model with individual differences being implemented as rates (far left, red lines) and a phase transition model based on differing transition times (middle left, blue lines). The right two panels combined simulated crash risk curves from the incremental and phase transition models (n=100 each). The far right panel shows that the aggregated (average) crash risk patterns are nearly identical, both resembling an incremental pattern.

Although the population-level crash patterns display an incremental decrease in crash rates, there is weak (if any) evidence of uniform individual-level *incremental* decreases in cognitions and behaviors associated with increased crash risk that match the power law

pattern in the population-level crash rate data. For example, research from naturalistic driving studies (i.e., studies with in-vehicle data recorders) and from survey-based cohort studies do not indicate that recently licensed drivers have uniformly and incrementally decreasing risky driving behavior that resembles a power law pattern (Carney, McGehee, Lee, Reyes, & Raby, 2010; McDonald, Sommers, & Fargo, 2014; Prato, Toledo, Lotan, & Taubman - Ben-Ari, 2010; Roman, Poulter, Barker, McKenna, & Rowe, 2015; Simons-Morton, Ehsani, Gershon, Klauer, & Dingus, 2017). Collectively, these data indicate that many new drivers are already driving relatively safely at the point of licensure and demonstrate fairly stable patterns of driving, while a minority of adolescents are more risky than their peers and more risky than experienced adults. Although these data provide important descriptive information, they cannot illuminate cognitive change processes. Thus, a new explanatory framework and a new methodological approach is needed to develop better theories of how novice drivers' crash risk reduces. These theories can then be used to make more effective interventions. In the following section we introduce the basic components of such a framework, a phase transition framework, which we believe provides a more parsimonious account of existing data.

Phase transitions have been observed across a variety of cognitive and behavioral domains (Lewis, 2000; Thelen & Ulrich, 1991). For example, strategic insight can lead to substantially improved performance in adult learners, and in their classic paper, Smith and Thelen reviewed how learning to walk can best be described as a series of phase changes within a dynamical systems framework (Gray & Lindstedt, 2017; Smith & Thelen, 2003). Applied to the context of transportation safety, the overall process of transitioning from a risky novice driver to a more experienced safer driver might be accomplished through a series of phase transitions, with abrupt transitions happening at different times for different drivers across different cognitive, emotional, social, and behaviors systems that support the driving task. This has direct implications for interventions that are designed to increase the incremental *accrual* of experience versus intervention approaches that facilitate achieving a phase change (i.e., a reorganization or qualitative change; this is known colloquially as a tipping point).

In summary, we propose three initial principles of a phase transition framework of individual-level, within-person changes in crash risk among recently licensed drivers: (1) changes in risk may be abrupt, non-linear, and occur at different times for different drivers. Phase states are generally stable, constrained by attractor dynamics (see Granic & Patterson, 2006; Smith & Thelen, 2003) and therefore, *meaningful* changes in risk are unlikely to be due to the incremental accumulation of pre-or post-license experience and are more likely to be attributed to constructs like strategy acquisition; (2) person-environment interactions can cause an immediate phase transition, or increase the probability of a future phase transition; and (3) there is substantial heterogeneity with respect to why any given novice driver might be at higher risk than his or her peers. In the current study, we provide the first direct test of the first two principles using a computational cognitive modeling approach (Sun, 2009). Detailed theoretical and methodological implications associated with applying a dynamical systems framework to develop a phase transition framework of driver behavior are explored in a separate paper (Mirman, Under Review).

2.0 General Methods

Using computational cognitive modeling we evaluated how well a phase transition model could approximate the population-level data from over 1 million adolescent drivers from four different jurisdictions across three different countries (Simulation 1) and compared this to an incremental (power-law) model predicated on experience accrual, which is the most commonly accepted theory about how drivers' crash risk reduction after licensure. Next, we modeled data from two US-based field trials of individual-level interventions administered prior to the accumulation of any post-licensure driving experience (Simulations 2 and 3). We chose these two experimental trials because they used rigorous experimental designs, used police-reported crashes as the primary outcomes, and they were administered prior to the accumulation of any post-license driving experience.

The purpose of computational cognitive modeling is to explicate the specific cognitive processes that underlie human learning and performance using formal (mathematical) algorithms. The collection of algorithms is what is called the computational "model". The model is then used to evaluate theoretical claims/assumptions and to generate novel predictions by conducting model simulations and evaluating the model results against observed data (for more discussion see Sun, 2008, 2009). This evaluation can be qualitative (the model produces the same general pattern as the observed data; e.g., a power-law-shaped "learning curve") or quantitative (the model produces a precise quantitative fit to the observed data; e.g., a close fit to actual crash rate data). In either case, the overall goal is to assess whether the theory that informed the model can explain the observed data. If the model cannot produce the observed behavioral data, that is a strong signal that the theory on which the model is built needs to be revised.

It is important to distinguish computational cognitive modeling from statistical modeling. A statistical model is a quantitative description of the observed data. A computational cognitive model is an implementation of the cognitive processes that are hypothesized to have generated the observed data. This is an important distinction: a linear regression makes no claim about *why* two variables have a linear relationship, it just describes the slope of that relationship; a computational cognitive model is designed to test a formal claim about *why* the two variables have a linear relationship. This kind of computational cognitive modeling approach is commonly employed in the fields of learning and cognitive science, but has not yet been applied to learning to drive. Although there is a strong tradition of computational modeling and simulation in the transportation sciences, these models have generally been restricted to predictive models of realtime action-motor control of experienced drivers and been informed by theories of experienced driver behavior.

Our approach is different in that it is the first attempt to model the cognitive changes that influence how novice drivers become, or change into, more experienced drivers. In the current paper we evaluate a phase transition model where new drivers' crash risk decreases relatively abruptly, with the timing and cause of the transitions varying across drivers. We chose a sigmoid function rather than a threshold function because an abrupt transition is not necessarily instantaneous and because the sigmoid provides a formalism for investigating effects of transition timing (the focus of this paper) as well as learning rate, thus providing a

more robust basis for future theoretical development. The simulation, analytic, and data visualization code is available via the Open Science Framework at https://osf.io/qbgpd/?view_only=8d6185e36e0f497fa548ece9fea38a28.

3.0 Simulation 1: Population-level Crash Rates

3.1 Methods

We mathematically defined the phase transition process with a sigmoid function over time, which produces an abrupt transition from a higher risk driver state to a lower risk driver state, simulated 10,000 drivers by randomly sampling phase transition times, aggregated the individual crash risk curves, and tested how well the aggregated pattern fit the normalized population crash rate data of recently licensed adolescent drivers from four different locales: New Jersey, USA (n=410,230), (Curry et al, 2015); North Carolina, USA (n=629,144), (Foss et al, 2011); Nova Scotia, Canada (n=40,661), (Mayhew et al, 2003) and Victoria, Australia (n not reported) (VicRoads, 2005). Note that we use crash risk to refer to a driver's learning-related risk (or psychological risk). When we simulate the crash rates there is another component that computationally captures other factors that contribute to a driver's risk for a crash, unrelated to how well they have learned to drive (e.g., impairment, no fault crashes, etc.).

The phase transition was defined as a sigmoid function [1], which produces a value above 0 (lowest possible risk state) and below 1 (highest possible risk state):

$$\text{Crash Risk} = \frac{1}{1 + e^{-g(-t+s)}} \quad (1)$$

Where t is time, initially in the range -3 to 3 to cover both the high and low plateaus of the sigmoid, then rescaled to correspond to months since licensure; g is the slope of the learning curve, set so that most of the transition happens within the span of one month, which is the temporal grain at which population-level crash data are typically reported and is a substantially faster decrease than would be predicted by an incremental model ($g = 15.0$ for these simulations); and s is phase transition shift – the relative time of onset of phase transition.

A population of 10,000 drivers was simulated by sampling the phase transition shift parameter, s , from an exponential distribution with a broad range and multiplied by the g value. These individual crash risk curves were aggregated and linearly scaled to fit the normalized population crash rate data. Following previous studies that used a power-law model to fit population-level crash rates (Foss et al, 2011), the crash rates were normalized so that the initial rate was 1.0 in all cases, thus putting them on the same scale and making the similarities in the pattern more apparent. These population crash rate curves were also fit using a power-law function [2], for comparison with previous studies: (Foss et al., 2011)

$$\text{Crash Risk} = a(\text{Month})^b \quad (2)$$

3.2 Results

Our first goal was to illustrate that the crash risk trajectories do not have to decrease incrementally following licensure by mathematically demonstrating that aggregating individual abrupt phase transition curves can produce equally good fits to incremental population-level crash rate curves (Table 1). The crash rate data (points), power law fits (red lines), and phase transition model fits (blue lines) for the four jurisdictions are shown in Figure 2.

In the next sets of simulations we address evidence from two very different individual-level interventions that were implemented prior to drivers having any on-road post-license experience and that produced effects that are larger than would be expected on an incremental model of crash risk reduction.

4.0 Simulation 2: On-road Driving Assessment (ODA)

Learner's permit holders (n=458) assigned to take a pre-license comprehensive on-road driving assessment (ODA) (Mirman et al., 2014) were an estimated 53% less likely to experience a police-reported crash during the post-license period than a comparison group (Mirman, Curry, Elliott, Long, & Pfeiffer, 2018). The ODA occurred approximately 12 and 24 weeks into the learner's permit period and was conducted by a certified driving rehabilitation specialist.

Unlike typical driver education and instruction administered during the learner period, the ODA lasted just under an hour, and intentionally exposed permit holders to a diverse array of challenging driving conditions and environments that required performance of prototypical driving behaviors inherent to each environment (e.g., changing lanes on a highway, merging in and out of traffic). The evaluators gave directional guidance, but did not educate adolescents during the drive. After the 24-week ODA was complete, the evaluators gave feedback on driving performance to parents and adolescents and answered questions (no feedback was provided at the 12-week ODA); see Mirman et al., 2018 for more detail.

An incremental change model of crash risk reduction associated with the accrual of post-license driving experience predicts that the ODA should have no meaningful effect on real-world crash rates because the ODA involves no post-license driving experience, lasted about an hour, and its observed effect was not an incremental one. It is important to note that the specific 53% estimate does not matter that much because under an incremental model the predicted effect should be negligible. If the ODA intervention shifts the distribution of drivers' phase transition timings so that more drivers have earlier phase transitions, then it could plausibly have a large effect on post-license crash risk.

4.1 Methods

The same sigmoid model of phase transitions was used for these simulations, but now with a more elaborated distribution of phase transition times. The phase transition shift parameter was sampled from a cumulative probability distribution that depended on (1) licensure status and (2) exposure to the ODA. That is, we mathematically implemented the assumptions that post-license independent driving increases the probability of a phase transition and that the

pre-license ODA experience also increases the probability of a phase transition. The first assumption (post-licensure experience is more likely to induce a phase transition than pre-licensure experience) mirrors the assumption that independent driving experience is more important for becoming a safer driver than pre-license experience (McCartt et al, 2009; Williams, Shope, & Foss, 2017). The second assumption (phase transitions are more likely after the ODA) reflects the fact that the ODA may have oriented people to their weaknesses or tolerances for task difficulty, which then encouraged them to try to figure out strategies to deal with those weaknesses and/or to make driving easier. That would take time and require opportunities to test out new strategies, which are likely to differ across individuals (as are weaknesses). In other words, the idea is that the ODA improved drivers' ability to self-calibrate. Thus, we modeled the ODA effect as increasing the probability that a phase transition would occur after the ODA, but not necessarily immediately upon ODA completion.

First, we defined cumulative probability distributions for phase transitions (left panel of Figure 3): For the first 6 months after study enrollment (participants had no more than 5 hours of driving practice at the point of enrollment), there was a slowly increasing proportion of participants that started their phase transition. Starting at 6 months post-enrollment (24 weeks), participants began to get licenses and the probability of phase transitions started to rise more rapidly (red line, corresponding to the Non-ODA group). The ODA group also had the ODA experience at 6 months and this produced an increase in probability of a phase transition (blue line, corresponding to the ODA group). These cumulative probability distributions were sampled randomly to create a sample of 5000 ODA and 5000 non-ODA phase transition curves (middle panel of Figure 3 shows the aggregates of those curves). The phase transition curves were then used to simulate crash events from a binomial distribution with probability proportional to the crash risk defined by the phase transition curve and a constant baseline crash risk [3]:

$$p(\text{crash}) = 0.013 (\text{Risk}) + 0.0013 \quad (3)$$

The crash risk component was suppressed during the early (pre-licensure) months on the assumption that the effect of a driver's inexperience is mitigated by other factors (e.g., the presence of their parent in the vehicle). The coefficients (0.013 and 0.0013) were selected to provide a reasonable qualitative fit to the behavioral data, but the specific values are not critical because model performance was consistent across a relatively large parameter range. Inherent to this approach is the assumption that the transition towards being a safer driver begins during the learner (pre-license) period and proceeds through the independent driving (post-license) period. Therefore, the model captures drivers' trajectories of decreasing risk and distinguishes the effect of experience by license status.

4.2 Results

The panels in Figure 3 proceed left-to-right according to the steps of the simulation. The left panel shows the cumulative probability distribution of phase transitions for the ODA (blue) and non-ODA (red) conditions. The middle panel shows the aggregate crash risk curves for

the two groups, each one corresponding to 5000 samples from the distributions in the left panel; drivers began to obtain licenses around month 6. These aggregate transition curves are not incremental; they have an initially shallower slope and sharp lower asymptote. Recall that the middle panel reflects *psychological* risk, and when it approaches zero this indicates that the driver has gotten about as close to floor as she or he is going to get with respect to how well she or he has learned to drive. There are other factors besides ability that affect crash risk and those are reflected in the intercept term in equation 3 (0.0013). The right panel shows the cumulative crash data from the ODA study (points) and the model predictions (lines) based on binomial sampling from the individual crash risk curves (whose aggregates are shown in the middle panel). Individual crashes were simulated as binary events with probabilities related to individual driver crash risk. The result was a good fit to the cumulative crash rate data from the ODA study (right panel).

5.0 Simulation 3: Risk Awareness and Perception Training (RAPT)

The Risk Awareness and Perception Training (RAPT) evaluation (n=5,251) was conducted in California state licensing offices. Recently licensed participants with no post-license driving experience were randomized to RAPT or a pre-test only condition. RAPT has been extensively tested in laboratory settings and is a computer-based training program primarily designed to improve drivers' visual search strategies, so that they are better able to identify emerging hazards (Fisher, Pollatsek, & Pradhan, 2006); minor modifications were made to adapt RAPT to be suitable for administration in licensing offices. Adolescents were shown photographs of traffic scenes and given two chances to use a cursor to indicate where they should be looking (e.g., at a particular hazard or area); if they failed to indicate the correct location, a red oval was used to indicate the hazard and brief written text was used to explain the rationale. Hazard anticipation training programs have been successful in improving visual search skills in simulator and controlled on-road evaluations (McDonald, Goodwin, Pradhan, Romoser, & Williams, 2015). The California evaluation was the first large-scale field trial of the RAPT program (Thomas, Rilea, Blomberg, & Korbela, 2016), and reported that the crash rates for male drivers who participated in RAPT were 23% lower than for male drivers who did not; no statistically significant effect was observed for female drivers (both the RAPT and comparison group females had equally low crash rates, comparable to the reduced crash rates of the RAPT group males). An incremental change model of crash risk reduction associated with the accrual of post-license driving experience predicts that RAPT should have no meaningful effect on real-world crash rates because the RAPT was very brief (less than 20 minutes) and involved no on-road experience.

5.1 Methods

To model the effect of the RAPT program on crash rates we used the same basic strategy as the simulations of the ODA effect. We first defined functions for cumulative probability of a phase transition occurring. As in the ODA simulations, we assume a lower rate of phase transitions prior to licensure and a higher rate of phase transitions after licensure. Because the RAPT program is explicitly designed to improve hazard anticipation by teaching hazard detection and recognition strategies, we assumed that the intervention induces an immediate, but modest (10%) increase in probability of a phase transition for participants who received

RAPT. Phase transition times were then sampled from these distributions for the comparison and RAPT groups ($n = 5000$ in each group). The phase transition curves were used to simulate crash events as a binary event with probability proportional to the crash risk defined by the phase transition curve and a constant baseline crash risk [4]:

$$p(\text{crash}) = 0.035(\text{Risk}) + 0.0028 \quad (4)$$

The slope and intercept for this function are somewhat higher than for the ODA simulation because the crash rates were higher in the RAPT study than in the ODA study; as in the ODA simulation, the coefficients (0.035 and 0.0028) were selected to provide a reasonable qualitative fit to the behavioral data, but the model performance was consistent across a relatively large parameter range.

5.2 Results

The results of the simulation are shown in Figure 4. The panels proceed left-to-right according to the steps of the simulation. The left panel shows the cumulative probability distribution of phase transitions for the RAPT (blue) and Comparison (red) conditions; month 0 corresponds to when an individual passed their on-road driver's license exam and when the RAPT intervention was administered. The middle panel shows the aggregate crash risk curves for the two groups, each one corresponding to 5000 samples from the distributions in the left panel. The right panel shows the crash data from the RAPT study (circles, vertical lines indicate standard errors) and the model predictions (diamonds) based on binomial sampling from the individual crash risk curves (whose aggregates are shown in the middle panel). This modest increase in phase transition probability matched the results of the RAPT trial (right panel).

6.0 Discussion

We demonstrated that a phase transition change framework can account for population-level crash patterns and the ability of brief pre-license interventions to reduce crash risk for novice drivers. On an incremental model of post-license crash risk reduction, the critical factor in decreasing crash risk is the accumulation of post-licensure driving experience (Foss et al., 2011; McCartt et al., 2009; Simons-Morton & Ehsani, 2016; Williams et al., 2017). In such terms, the crash rate reduction associated with the ODA pre-license experience is equivalent to approximately 18 months of post-license driving experience and crash rate reduction associated with the RAPT experience is equivalent to approximately 3-6 months of post-license driving experience. If trajectories in novice drivers' motor vehicle-related crash risk are characterized by incremental processes, undergirded primarily by incremental accumulation of post-license driving experience, then neither of these interventions should have had a meaningful effect on crash involvement. We also demonstrated that the commonly observed population-level incremental decrease in crash rates can easily be generated by aggregating asynchronous abrupt changes in individual-level crash risk – thus causing a power-law artifact in the aggregate crash data. These results indicate why it is problematic to make inferences about the nature of how individual driver's crash risk reduces using patterns in aggregate crash data; two very different individual-level theories –

phase transition and incremental – are capable of generating the same pattern in the aggregated crash data.

6.1 Implications for Reporting and Interpreting Research Results

Implicitly relying on the incremental model may lead to unintentionally making misleading statements about how novice drivers' risk changes with experience. Common examples include referring to exposure-based changes in the aggregate crash data as “learning rates” or interpreting them as such, and then calculating the average risk reduction for any given driver using the population-level data; so for example saying that an average driver’s risk will reduce by Y amount after X months of licensing or X miles driven.

An arithmetic average can always be calculated, but the notion of an “average driver” is deeply problematic in the context of aggregation artifacts such as the one demonstrated in Simulation 1. The phrase “average driver” implies that such a driver actually exists and that it is typical or representative of other drivers. In the Phase Transition Model there may not be any such driver, nor any drivers that resemble this hypothetical average driver.

Additionally, a population cannot learn; learning in the psychological sense happens only at the individual-level. In other words, “learning” is not an attribute that a population can have.

Crucially, it is not logically valid to assume that individual-level changes in crash risk mirror the population-level changes in crash rates. In statistics, this is known as an ecological fallacy and in formal logic it is known as the fallacy of division, a type of category error. We want to stress that the traditional interpretation of the aggregate crash data as reflective of an incremental change pattern is a very intuitive and statistically accurate interpretation of the aggregate crash data. At stake here is becoming aware of implicit and important assumptions about what the aggregate pattern means for *individual* drivers. Belief biases exert a strong influence on reasoning and can lead to accepting arguments that are not logical when the conclusion is believable or otherwise appealing. Table 2 shows three example arguments with identical logical structure, though the conclusion of example A probably seems reasonable whereas examples B and C seem much less so.

Most people will intuitively understand why it does not make sense to generalize the aggregate pattern to individual coin flips or conception (Table 2). However, in the case of learning to drive, the dominant intuition among transportation safety scientists has been that the incremental decrease in the aggregate crash rate reflects incremental individual-level decreases in crash risk (see McCartt et al., 2009). Thus, we recommend not using the phrase “learning rate” to refer to population-changes in aggregate crash patterns because it can cause readers to make the fallacy of division when interpreting study results. Instead of learning rate, authors may want to consider using phrases such as “change in population-level rate” to make it clear that the phenomena is being described at the population-level and not at the individual-level. Finally, it is possible that the individual-level change patterns of different driving phenomena *can* match the population-level change pattern, but it cannot be assumed.

6.2 Implications for Interventions

Driver training and behavior change programs are implemented on the individual-level and not the population-level. It is critical to know what the population-level pattern looks like, but, when we create individual-level interventions that are primarily informed by change patterns in the population-level crash data we run a high risk of creating intervention approaches that are based on incorrect assumptions about the psychological reasons why a driver's risk might be changing. This point becomes clearer if we think about interventions designed to target the "average driver" using the example from the prior section. Recall that this average hypothetical driver may not actually exist in the target population, thus it is not likely that an intervention designed specifically for that driver (or type of driver) will be effective.

For example, McCartt et al., (2009) estimated that during the first year of driving seventeen year-old drivers will see their risk reduce by 30% due to "experience" (pg 215). On an incremental model, this would be interpreted as meaning that at least some seventeen year-old drivers see a 30% risk reduction and many seventeen year-old drivers see a reduction somewhere near 30%. Alternatively, in an extreme example of the phase transition model, a possible interpretation is that half of the drivers would see a 60% reduction and the other half would see no reduction at all, while no individual seventeen year-old driver would see a 30% reduction or anything near that value. From an intervention perspective, the goal is then to identify the factors that put someone into the risk-reduction group rather than the no-reduction group and to shift more drivers into the risk-reduction group.

The most successful approach for mitigating novice drivers' crash risk includes limiting their exposure to contexts known to place them at higher risk (e.g., driving at night and with same-aged passengers) (Williams, 2017). Historically, individual-level interventions intended to make novice drivers safer and more competent prior to licensure (e.g., driver education) have been largely unsuccessful at reducing crash risk, and at times even harmful (Beanland, Goode, Salmon, & Lenné, 2013; Christie, 2001; Mayhew et al., 2003). Similarly, with a few exceptions there is not strong evidence that other types of individual-level behavioral interventions and training programs are effective at reducing adolescents' crash risk (Curry, Peek-Asa, Hamann, & Mirman, 2015; McDonald, Goodwin, Pradhan, Romoser, & Williams, 2015). A poor understanding of how crash risk decreases has made it difficult to develop effective individual-level interventions.

Our results suggest that researchers should consider cognitive change processes that beget translational transformations in crash risk trajectories (i.e., shifting crash risk curves) instead of thinking primarily about learning rates and the relationship between *accruing* experience and incremental changes in crash reduction. Strategy acquisition is one such process. Learning a new strategy can produce an abrupt improvement in performance (i.e., a leap) and may not require direct experience if the strategy can be learned through other means. Indeed, without explicitly referencing the phase transition framework, the design of the RAPT program assumed that explicit training in a virtual (computer-based) environment can allow individuals to quickly learn hazard detection strategies that might transfer to safer real-world driving. Our simulations show that if even a small subset of the participants benefitted from the RAPT program, then there could be a substantial effect on subsequent crash rates.

Inadequate surveillance was found to be the critical reason for about 20% of the errors made by novice drivers involved in severe MVCs (Curry, Hafetz, Kallan, Winston, & Durbin, 2011). RAPT was explicitly designed to improve drivers' visual search strategies. Improved visual search strategies can facilitate hazard recognition and lengthen the amount of time adolescents have to avoid a collision.

The mechanisms behind the ODA are less straightforward as it was a more complex experience designed to simulate on-road post-license driving during the pre-license period. The ODA included exposing the novice driver to challenging, diverse, live traffic, without instructional support by an adult co-driver. The ODA also entailed giving feedback to the drivers and their parents about how to improve after the assessment was completed. It is possible that the ODA experience rapidly sensitized adolescents to their weaknesses in the face of the true task demands of real-world driving, which are usually reduced by driving supervisors during the supervised driving period. Perhaps knowing more about their weaknesses encouraged adolescents to develop strategies that reduced the task demands of driving, like traveling at slower speeds in certain circumstances (e.g., entering a curve, when sight lines are obstructed) and taking steps to reduce cognitive load (e.g., avoiding multi-tasking). Decision errors (e.g., going too fast for road conditions, misjudging a curve) account for 40% of the critical errors that were observed in adolescents' MVCs (Curry et al., 2011). Almost all of these decisions are made under uncertainty (i.e., with incomplete information) in time-sensitive scenarios with a clear benefit for fast and accurate decision-making.

Our modeling approach reflected a possible key difference between how the ODA and RAPT may have affected phase transition timing: inducing immediate phase transitions (RAPT) and increasing the probability of a phase transition (ODA). This distinction between inducing immediate phase transitions and increasing the probability of spontaneous phase transitions is not critical for accounting for the RAPT and ODA data – the model is quite robust and produces largely the same patterns over a range of parameters and implementations – but it illustrates different ways that hypothesized effects of interventions can be concretely (i.e., mathematically) connected to the phase transition framework. More research on this topic is important.

6.3 Theoretical Implications

Using computational cognitive modeling methods we introduced a new phase transition framework to guide research on real-world adolescent risk-taking behavior indexed by motor vehicle crashes that indicates evidence for strategy acquisition processes as underlying changes in adolescents' crash risk. We do not argue that the journey from novice to expert could be explained through a single phase transition; rather there may be a series of transitions that relate to different aspects of driving. These abrupt improvements could be interspersed with periods of incremental improvement; for example, discovery of a new superior strategy (abrupt reduction in risk), followed by increasing success in correctly using that strategy (additional incremental reductions in risk) (Gray & Lindstedt, 2017).

Further, the phase transition framework does not presuppose that “learning to drive” in its broadest sense is a scripted process that all young drivers must pass through uniformly.

Practically, these ideas mean that crash risk may drop abruptly for different reasons, at different times, for different adolescent drivers. A large reduction in crash risk for one teen driver might be precipitated by completely different factors than another teen driver. Theories about how novice drivers' risk reduces must be based on a model of psychological systems rather than on a statistical description of population-level crash rates.

We see the phase transition theory of change and its implementation in the computational cognitive model as a starting point from which more specific models of the processes that give rise to phase transitions can be built. For example, in our models of the RAPT and ODA effects we implemented only one phase transition, but we expect that becoming a safer driver involves multiple phase transitions across different time scales and domains. To guide future research efforts we recommend utilizing coordinated within-person study designs, in the real-world and in laboratory settings, which permit the detection of phase transitions and their precursors (Mirman, Under Review).

Our generative model accounted for data from large epidemiological studies (over a million adolescent drivers from multiple jurisdictions) and behavioral studies with strong experimental designs and high ecological validity. The original trial reports of the RAPT and ODA studies describe their limitations, ranging from concerns about sample selection bias (e.g., safe samples interested in safety) to difficulty of detecting effects on crashes because they are rare events. (Mirman et al., 2018; Thomas et al., 2016). Additional behavioral studies are critical, but these two studies represent foundational starting points for two important reasons: (1) they used much larger sample sizes than the typical small laboratory-based samples of adolescents and (2) they utilized police-reported motor vehicle crashes—a critically important, ecologically valid, and objective outcome variable—in the context of an experimental design.

6.4 Conclusion

Cross-level (individual $\longleftrightarrow$ group) assumptions must be made explicit and tested because the aggregated data are underdetermined with respect to the kinds of change processes that are underway at the individual-level. In this paper we demonstrated that an incremental decrease in aggregate crash rate can reflect both incremental reductions in crash risk within individuals and an incremental increase in the proportion of drivers who have experienced an abrupt decrease in crash risk. Therefore, while it is true to say that the population of novice drivers' crash risk reduces in the months following licensure, it is not necessarily true to say that an individual young driver's crash risk reduces in the months following licensure. Further, the phase transition framework provides a parsimonious account of how two brief interventions produced substantial decreases in crash rates that are far larger than can be explained by incremental learning. In sum, compared to the incremental model, the phase transition framework is: (a) much more consistent with other literature on human development, learning, and performance and (b) offers a much more parsimonious way to explain data at both the individual level and the population level than the incremental model, which has been implicitly adopted by the transportation safety field.

References

- Carney C, McGehee DV, Lee JD, Reyes ML, & Raby M (2010). Using an event-triggered video intervention system to expand the supervised learning of newly licensed adolescent drivers. *American Journal of Public Health*, 100(6), 1101–1106. 10.2105/AJPH.2009.165829 [PubMed: 20395588]
- Curry AE, Hafetz J, Kallan MJ, Winston FK, & Durbin DR (2011). Prevalence of teen driver errors leading to serious motor vehicle crashes. *Accident Analysis and Prevention*, 43(4), 1285–1290. 10.1016/j.aap.2010.10.019 [PubMed: 21545856]
- Curry AE, Pfeiffer MR, Durbin DR, & Elliott MR (2015). Young driver crash rates by licensing age, driving experience, and license phase. *Accident Analysis & Prevention*, 80, 243–250. 10.1016/j.aap.2015.04.019 [PubMed: 25939133]
- Fisher DL, Pollatsek A, & Pradhan AK (2006). Can novice drivers be trained to scan for information that will reduce their likelihood of a crash? *Injury Prevention*, 12 Suppl 1(6 2009), i25–i29. 10.1136/ip.2006.012021 [PubMed: 16788108]
- Foss RD, Martell CA, Goodwin AH, & O'Brien NP (2011). *Measuring Changes in Teenage Driver Crash Characteristics During the Early Months of Driving*. Washington, DC: AAA Foundation for Traffic Safety 10.1037/e554022012-001
- Granic I, & Patterson GR (2006). Toward a comprehensive model of antisocial development: a dynamic systems approach. *Psychological Review*, 113(1), 101–131. 10.1037/0033-295X.113.1.101 [PubMed: 16478303]
- Gray WD, & Lindstedt JK (2017). Plateaus, Dips, and Leaps: Where to Look for Inventions and Discoveries During Skilled Performance. *Cognitive Science*, 41(7), 1838–1870. 10.1111/cogs.12412 [PubMed: 27766673]
- Haider H, & Frensch PA (2002). Why aggregated learning follows the power law of practice when individual learning does not: comment on Rickard (1997, 1999), Delaney et al. (1998), and Palmeri (1999). *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 28(2), 392–406. 10.1037/0278-7393.28.2.392
- Lewis MD (2000). The Promise of Dynamic Systems Approaches for an Integrated Account of Human Development. *Child Development*, 71(1), 36–43. 10.1111/1467-8624.00116 [PubMed: 10836556]
- Mayhew DR, Simpson HM, & Pak A (2003). Changes in collision rates among novice drivers during the first months of driving. *Accident Analysis and Prevention*, 35(5), 683–691. 10.1016/S0001-4575(02)00047-7 [PubMed: 12850069]
- McCartt AT, Mayhew DR, Braitman KA, Ferguson SA, & Simpson HM (2009). Effects of age and experience on young driver crashes: Review of recent literature. *Traffic Injury Prevention*, 10(3), 209–219. 10.1080/15389580802677807 [PubMed: 19452361]
- McDonald CC, Goodwin AH, Pradhan AK, Romoser MRE, & Williams AF (2015). A Review of Hazard Anticipation Training Programs for Young Drivers. *The Journal of Adolescent Health*, 57(1 Suppl), S15–23. 10.1016/j.jadohealth.2015.02.013 [PubMed: 26112734]
- McDonald CC, Sommers MS, & Fargo JD (2014). Risky driving, mental health, and health-compromising behaviours: risk clustering in late adolescents and adults. *Injury Prevention*, 1–8. 10.1136/injuryprev-2014-041150 [PubMed: 23728529]
- Mirman JH, Curry AE, Elliott MR, Long L, & Pfeiffer MR (2018). Can Adolescent Drivers' Motor Vehicle Crash Risk Be Reduced by Pre-Licensure Intervention? *Journal of Adolescent Health*, 62(3), 341–348. 10.1016/j.jadohealth.2017.09.015 [PubMed: 29223562]
- Mirman JH, Curry AE, Schultheis MT, Brant C, Wang W, Winston FK, & Durbin DR (2014). Development of On-Road Driving Assessment for Learner Teen Drivers. *Transportation Research Record*, 2465, 64–72. 10.3141/2465-09
- Mirman JH (Under Review). A dynamical systems perspective on driver behavior.
- Murre JMJ, & Chessa AG (2011). Power laws from individual differences in learning and forgetting: mathematical analyses. *Psychonomic Bulletin & Review*, 18(3), 592–7. 10.3758/s13423-011-0076-y [PubMed: 21468774]

- Myung IJ, Kim C, & Pitt MA (2000). Toward an explanation of the power law artifact: insights from response surface analysis. *Memory & Cognition*, 28(5), 832–840. 10.3758/BF03198418 [PubMed: 10983457]
- Prato CG, Toledo T, Lotan T, & Taubman - Ben-Ari O (2010). Modeling the behavior of novice young drivers during the first year after licensure. *Accident Analysis and Prevention*, 42(2), 480–486. 10.1016/j.aap.2009.09.011 [PubMed: 20159070]
- Roman GD, Poulter D, Barker E, McKenna FP, & Rowe R (2015). Novice drivers' individual trajectories of driver behavior over the first three years of driving. *Accident Analysis and Prevention*, 82, 61–69. 10.1016/j.aap.2015.05.012 [PubMed: 26047833]
- Seigler RS (2006). Microgenetic analysis of learning In Kuhn D, Seigler RS, Damon W, & Lerner RM (Eds.), *Handbook of Child Psychology* (6th ed., pp. 464–510). Wiley.
- Simons-Morton BG, & Ehsani JP (2016). Learning to drive safely: reasonable expectations and future directions for the learner period. *Safety*, 2, 1–19. 10.3390/safety2040020
- Simons-Morton BG, Ehsani JP, Gershon P, Klauer S, & Dingus T (2017). Teen Driving Risk and Prevention: Naturalistic Driving Research Contributions and Challenges. *Safety*, 3(4), 29 10.3390/safety3040029
- Smith LB, & Thelen E (2003). Development as a dynamic system. *Trends in Cognitive Sciences*, 7(8), 343–348. 10.1016/S1364-6613(03)00156-6 [PubMed: 12907229]
- Sun R (2009). Theoretical status of computational cognitive modeling. *Cognitive Systems Research*, 10(2), 124–140. 10.1016/j.cogsys.2008.07.002
- The Science of Adolescent Risk-Taking: Workshop Report. Committee on the Science of Adolescence. (2011). Washington, D.C.
- Thelen E, & Ulrich BD (1991). Hidden Skills: A Dynamic Systems Analysis of Treadmill Stepping during the First Year. *Monographs of the Society for Research in Child Development*, 56(1), 1–98. 10.2307/1166099
- Thomas FD, Rilea SL, Blomberg RD, C., P. R., & Korbela KT (2016). Evaluation of the safety benefits of the risk awareness and perception training program for novice teen drivers. Report No. DOT HS 812 235. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/1591962>
- VicRoads. (2005). *Young Driver Safety and Graduated Driver Licensing*. Kew.
- Williams AF, Shope JT, & Foss RD (2017). What is learned in becoming a competent driver? In Fisher DL, Caird JK, Horrey WJ, & Trick LM (Eds.), *Handbook of Teen and Novice Drivers* (1st ed., pp. 453–456). Boca Raton, FL.

Highlights

- Frameworks that account for changes in novice drivers' crash risk are underdeveloped.
- Changes in novice drivers' crash risk trajectories can be abrupt, not incremental.
- Abrupt changes in risk can signal a phase transition.
- Phase transitions may be caused by strategic learning processes.
- Phase transitions can be induced through intervention.

Figure 1.
Incremental and phase transition (discontinuous) models of drivers' changing crash risk aggregate into similar patterns.

Figure 2.

Normalized crash rates over the first two years post-licensure from four jurisdictions.

Points show the observed crash rates, red lines show the power law model fits, and blue lines show the phase transition model fits.

Figure 3.
Phase transition model of the effect of the ODA on post-licensure crash rates.

Figure 4.
Phase transition model of the effect of the RAPT program on post-licensure crash rates.

Table 1.

Goodness of fit (quasi- R^2) for incremental-accrual and phase transition models of population-level crash rates.

Jurisdiction	Power Law	Phase Transition
North Carolina, USA	0.955	0.977
Victoria, Australia	0.959	0.979
Nova Scotia, Canada	0.941	0.934
New Jersey, USA	0.930	0.962

Table 2.

The logical validity of an argument is independent of the believability of the conclusion

	Logical Argument	A Incremental-Accrual Model	B Pregnancy Accrual	C Coin Flips
Premise 1	Whole A (aggregate) has properties X and Y	Drivers' aggregate crash rate incrementally decreases with the accrual of experience.	Pregnancy rates among women incrementally increase from age 20 to 35.	The aggregate average of coin flips is 50% heads.
Premise 2	P is part of A	Individual drivers are part of the aggregate crash rate.	Individual women are part of the aggregate pregnancy rate.	Individual coins are part of the aggregate heads rate.
Conclusion 1	P has properties X and Y	Individual drivers' crash risk decreases incrementally with the accrual of experience.	An individual woman's pregnancy status incrementally increases from age 20 to 35.	It is possible for an individual coin flip to be 50% heads.